

UNITED STATES PATENT APPLICATION

FOR

**METHOD AND SYSTEM FOR BUILDING A REPORT FOR
EXECUTION AGAINST A DATA STORE**

Inventor(s): **Aaron Stephen Meyers**, a citizen of the United States, residing at
2018 West 1600 North
Provo, UT 84604

Robert Alan Meyers, a citizen of the United States, residing at
1762 South 100 East
Orem, Utah 84058

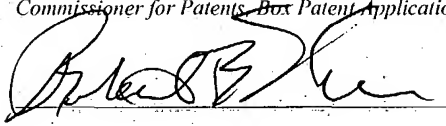
Eric Stephen Meyers, a citizen of the United States, residing at
2018 West 1600 North
Provo, UT 84604

Assignee: **ActiveViews, Inc.**
2018 West 1600 North
Provo, UT 84604

"Express-Mail" Label Number EU432952851US

Date of Deposit July 25, 2003

I hereby certify that this paper or fee is being deposited with the United States
Postal Service "Express Mail Post Office to Addressee" service under 37 CFR
1.10 on the date indicated above and is addressed to the Assistant
Commissioner for Patents, Box Patent Application, Washington, D.C. 20231.


Date 7/25/03

METHOD AND SYSTEM FOR BUILDING A REPORT FOR EXECUTION AGAINST A DATA STORE

5

Inventors:

Aaron Stephen Meyers
Robert Alan Meyers
Eric Stephen Meyers

10

BACKGROUND OF THE INVENTION

The present invention relates to information processing and more particularly
15 to query and reporting systems and methods related to information processing.

Organizations have always been large collectors of data. During the past two
decades, the sharply increasing popularity and advancement of personal computers have
introduced many new options and compelling opportunities for amassing great stores of data.
Notwithstanding, much of the value in this captured data lies in the ability to retrieve and
20 review it in certain ad hoc selections and presentations in order to discover information
contained in the data.

Historically, institutions have generated countless reports and extracts of data
to alleviate this problem. In most instances, programmers who understand the systems used
to capture data develop these reports and extracts. However, developing suitable extraction
25 and reporting tools for generating ad hoc queries against available data frequently has
required too much time and human capital, significantly limiting the usefulness of the data.

Legacy systems have significantly contributed to the problem. In the 1970s, mainframe computers comprised the center of most database systems development. The 1980s brought new mini-computer platforms. The late eighties and early nineties brought the client-server architecture running on personal computer hardware and popular server platforms such as Unix, NetWare and Windows. Despite these changes in platforms, architectures, tools, and technologies, large amounts of data and large numbers of data analysis applications continue to reside in the mainframes of the 1970s. By some estimates, more than 70 percent of existing institutional data still resides on mainframes.

To cope with this problem, several significant technological advances have been made. First, the advent of more powerful personal computers and operating systems has resulted in an explosion of data storage in standardized data stores or databases and desktop tools to manipulate this data. Second, technological advances have resulted in newer enterprise-class applications to manage data and an understanding of the fundamental differences between transactional and analytical systems. These advances have given rise to the “data warehouse” concept for consolidating resources around data retrieval. Third, display applications have made significant progress in providing users with easier and better access to data.

Although the foregoing technological advances have made significant progress towards solving the problems associated with extracting useful information from data, several problems still exist. To be truly useful, information needs to be accessible and useable by a broad cross-section of information consumers. Current systems are too complex, too expensive, too rigid and too insecure to support broad access and use of information.

Prior systems are too complex to support access and use by a broad cross-section of information consumers. The bulk of the world's population simply does not have the skills necessary to use them. Most will likely never have the skills necessary to compose and submit database queries. Some have suggested that desktop applications are an answer to this problem. However, desktop database solutions typically result in fragmented data that is oriented towards very specific needs. Moreover, this approach to data management assumes the end user has the time to expend on managing the data in spreadsheets, files, and desktop databases. While some users may be proficient at data management, most undertake these tasks as a necessity. Given the choice, most users find it more efficient to focus on actual analysis.

Moreover, the solutions proffered by prior systems to address this complexity often do not work well. Many require information users to work through experts. However, many information consumers cannot provide specific, accurate requirements to a report writing expert on the first try. Thus, a typical query/report cycle involves several iterations to get the data and presentation correct and to validate the results.

Typically, prior systems offer few tools, if any, to unsophisticated information consumers for quickly and easily creating reports without programming assistance. Existing tools are limited to simplified wizard and web interfaces that simplify complexity by minimizing the number of options and inputs required of users. However, because designers of these wizards and web interfaces presume that specific types of data will be manipulated, these interfaces may be inflexible and lead to inaccurate tables or cross sections of data when

applied to different types of data. They may also restrict the information user from accessing needed information.

Prior systems have often been too expensive to support broad access and use.

The emergence of special data warehouses deploying online analytical processing (OLAP)

5 and complex data models requires that data be transformed into multidimensional arrays or cubes, limiting both flexibility and timeliness of the data to some extent and requiring

expensive applications and programming support. Because of such, current systems are so

complex that they can only be deployed successfully in large enterprises that can afford

them. Millions of small-to-midsize institutions also need access to business data for

10 operational and decision-making purposes. While these institutions may not have formal, specialized data warehouses, they do need to perform reporting and analysis functions.

Regardless of an organization's size, improving the ability to access and use data is critical to success.

Prior systems have frequently been too rigid to support broad access and use.

15 Current systems limit user interaction with data. Users run reports typically created by experts. Users cannot generally look at fields or relations other than those included by experts. Users generally cannot drill down on data or drill through data to determine if the data makes sense.

Moreover, the various physical data layouts used in analytical systems have

20 diminished the value of the data for many information users. Most transactional systems employ an entity-relationship (ER) data model. ER is a modeling technique that seeks to remove redundancy in data, thus increasing transactional speed. ER removes redundancy by

recognizing that pieces or fields of information are related to other fields of information through one-to-one, one-to-many, and many-to-many relationships and by organizing these fields in several interrelated tables. A byproduct of ER modeling is that the relations incorporated into the data model generally embody business or logical relations. These relations are inherently understandable by humans because they are based upon the way humans manage data. Humans inherently know that customers have names, address and phone numbers, and that they place orders. Humans inherently know that orders are shipped through shippers, occur in sales regions, are taken by specific sales personnel and are for certain products.

Because ER-modeled systems result in numerous database tables, analytical systems typically do not handle them very well. As described above, transactional systems are generally modeled for data integrity and transactional speed. Typically, transactional systems only need to handle relatively small amounts of information. Most transactional systems only maintain data for a day, week, month or year. Accordingly, and because they typically utilize an ER data model, retrieving data from a transactional system often requires joining several database tables. Because transactional databases are relatively small, these joins do not require extensive processing power or memory. Analytical systems, however, are built to analyze large quantities of data. Accordingly, analytical systems contain a great deal of data. Frequently, data in analytical systems span several years or the entire life cycle of the associated institution. Because they contain extensive data, executing joins on multiple tables typically requires a great deal of processing power and memory.

Analytical systems have attempted to overcome excessive processing and memory needs through alternative physical data modeling strategies. Most of the existing data modeling strategies require transformations of transactional data from their ER-modeled structures into more analytically friendly, flat data models. While these efforts are designed to increase processing speeds, they often result in the loss of some data and important business logic.

Because entity-relations are generally not maintained by analytical systems, existing ad hoc data query tools present flat data layouts to users. Typically, a user wishing to create a query against an analytical system is presented with a myriad of fields from which to choose. Sometimes, these fields are grouped into logical groupings. However, in each case the inherent business logic relations among the fields are typically lost. Further, once the user has selected certain fields, existing data query systems either automatically decide what joins to make on the tables associated with the fields, or present the user with a clumsy Boolean interface that itself is not aware of the relations among the selected fields. This can produce inconsistent results as fields are added or removed from a query, and makes generating a meaningful query very difficult. Missing or automatically interpreted join information can cause cross joins to be performed, among other errors, some of which are detected by the query reporting engine, while others are not, transparently returning erroneous results to the user and making report validity and data integrity a major issue for an organization.

Prior systems are also typically too insecure to support broad use. Most organizations cannot provide broad access to analytical data because of a lack of security for the data. Current systems rely on the underlying database system to provide security for the

data. However, most database systems enforce security on an entire database or on tables or rows of the database. Therefore, organizations wanting to provide broad access to data are either required to expose sensitive data or restrict data access.

Thus, there is a need for a data query and reporting system that (1) allows
5 non-technical computer users to build complex queries, (2) minimizes the need to have technical computer users build complex database queries by hand; (3) accommodates the need for a simple, easy-to-understand iterative system for generating and validating queries; (4) reduces the complexity and costs associated with retrieving useful information for organizations of all sizes, especially smaller businesses; (5) minimizes data transformations
10 from transactional to analytical systems, preserving data and business logic, and promoting smarter queries; and (6) supports a robust security model, enabling organizations to protect sensitive data while providing broad data access.

SUMMARY OF THE INVENTION

According to the present invention, a system and method for building queries for execution against a data store is provided. The system and method empowers novice or casual computer users to easily and iteratively select elements or fields of the data store to
5 create useful reports.

According to one embodiment, the present invention is implemented through a distributed application that runs on multiple computers but is displayed on a graphical user interface (GUI). This GUI, combined with common input devices such as a mouse and keyboard, minimizes the learning curve for use of the present invention. Thus, even a novice
10 or casual user may quickly and easily understand and apply the present invention to a data store.

The present invention provides a simple-to-use data query and reporting system that retrieves data from a data store according to a user's desires in response to simple and efficient input commands. Using the invention, a user may select fields and relations
15 associated with data. A user may also group data records by row, column, or by row and column. The user may also summarize the contents of select numeric or aggregation fields and present the results in the same report. Summary results may be presented for an entire report or for logical groupings of data.

The present invention defines a data store in terms of a relational abstraction.
20 The relational abstraction generally parallels the entity-relationship inherent in transactional relational database management systems. Doing so preserves the business logic associated with such transactional systems for use by users of the invention. One skilled in the art will

readily recognize that an entity-relationship abstraction may also be applied to data storage systems that are not in the genre of traditional relational database management systems.

The data store definition describes views, fields and relations of the data store. View definitions identify tabular structures of rows and columns in the data store. Field definitions describe columns of data accessible to the user in a particular view. Relation definitions describe associations between various views. Typically such definitions are associated with one or more tables and columns of a conventional relational database management system. However, one skilled in the art will recognize that any means of providing an entity-relationship view on data may be used as part of the invention.

The power of the present invention comes from empowering users to select a single base view as the starting point for generating a query. This base view is used to constrain the selection and creation of report fields.

The base view is used in the selection and creation of report fields as the starting point of a relation path to a destination view of a report field. Each relation path contains zero or more relationships and is dynamically generated by the invention as a report is built. In building a report, fields are selected or relationships are followed iteratively and recursively. In this fashion, the present invention enforces the cardinality of relations relative to the base view.

If a relation path is empty or includes only to-one relationships, a selected field is considered scalar relative to a row in the base view. If a relation path includes at least one to-many relationship, a selected field is considered aggregate relative to the base view,

and must be paired with a function or expression that aggregates its values into a scalar value for each row in the base view.

The present invention improves report filtering and provides a means for quickly and easily constraining data based upon the base view of a report. In similar manner
5 to iteratively and recursively selecting scalar and aggregate fields, fields may be selected as filter fields. As filter fields are selected, logical filter operators may be applied, filter fields may be nested into filter groups, and filter group operators may be applied.

As a user selects fields associated with the base view, an embodiment of the present invention iteratively and recursively captures user input for a report, and
10 automatically generates an SQL query, which upon execution returns the desired result set of data. The SQL query includes a FROM clause based upon the base view's underlying SQL table or view. Fields on the base view with no relation path are selected and filtered directly in the query. Fields with a relation path cause nested subqueries to be generated, which are joined to the view of the outer query using the join keys specified for the relation.

15 One nested subquery is generated for each relation in the path, each containing a FROM clause for the relation's destination view. The subquery for each level of the relation path is nested within the previous level, similar to the recursive tree structure. Fields with fully or partially matching relation paths may reuse the same nested subqueries. The nested subqueries allow data across to-many relations to be aggregated before being
20 joined with the base view rows. This approach allows reports involving complex aggregation to be returned using a single SQL query, without usage of temporary tables or intermediate

processing. The present invention thus provides a means of easily validating reports, as the captured user request closely mirrors the automatically generated SQL query.

According to one embodiment, the present invention provides a robust security model, enabling organizations to protect sensitive data while providing broad data access. According to this embodiment, a person or process knowledgeable about the data contained in a data store and the organizational requirements to protect the data defines one or more security principals who will be given access to information contained in the data. Security principals may be users or groups of users, and the identities of and other information about security principals are typically maintained in a protected file, database or directory.

According to this embodiment of the present invention, the person or process also creates an entry in an access control list for each element of the data, specifying in the entry if the security principal is to be granted or denied access to the element. The data elements may be course grained, as in the case of entire tables or views, or fine grained, defining permissions for specific fields, ranges of fields, or field aggregates.

According to this embodiment,, the person or process also creates a security filter for each report view to be protected. The security filter must include at least one row of the view, otherwise data from the view could not be viewed. The security filter also includes an access control entry denying access to one or more specific data elements that might be specified in the view. Defining the access control entry to deny access permits the security filter to operate in a fashion similar to other filters pertaining to the invention.

Once security principals, access control entries and security filters are defined, access control is enforced on a report by requiring that each security principal be authenticated and by applying the security filters associated with the report according to the permissions associated with the security principal as defined in the access control list and the security filter.

That the invention improves over the drawbacks of prior database query and report applications and accomplishes the advantages described above will become apparent from the following detailed description of preferred embodiments and the appended drawings and claims.

10 BRIEF DESCRIPTION OF THE DRAWINGS

Other features and advantages of the present invention will be apparent from the following Detailed Description taken in conjunction with the accompanying Drawings, in which:

FIGURE 1 is a block diagram of a distributing computing system that provides an exemplary operating environment for the present invention.

FIGURE 2 is a tabular diagram of a sample database.

FIGURE 3A is a tabular diagram of certain metadata software objects associated with the sample database depicted in FIG. 2.

FIGURE 3B is a tabular diagram of certain metadata software objects associated with the sample database depicted in FIG. 2.

FIGURE 3C is a tabular diagram of certain metadata software objects associated with the sample database depicted in FIG. 2.

FIGURE 3D is a tabular diagram of certain metadata properties associated with an embodiment of the present invention.

5 FIGURE 3E is a series of tables illustrating eXtensible Markup Language (XML) examples of metadata software objects associated with an embodiment of the present invention.

FIGURE 4 is a main display window of an embodiment of the present invention.

10 FIGURE 5 is a window display illustrating a software wizard used in an embodiment of the present invention.

FIGURE 6A is a window display illustrating selection of a database according to an embodiment of the present invention.

15 FIGURE 6B is a window display illustrating selection of a base view according to an embodiment of the present invention.

FIGURE 7 is a window display illustrating a detail field drop area according to an embodiment of the present invention.

FIGURE 8 is a window display illustrating a group field drop area according to an embodiment of the present invention.

FIGURE 9 is a window display illustrating a measure field drop area according to an embodiment of the present invention.

FIGURE 10A is a logic flow diagram illustrating a method for generating a report according to an embodiment of the present invention.

5 FIGURE 10B is a logic flow diagram illustrating a method for displaying and creating queries according to an embodiment of the present invention.

FIGURE 11A is a window display and related XML snippet illustrating iterative generation of a report definition according to an embodiment of the present invention.

10 FIGURE 11B is a window display and related XML snippet illustrating iterative addition of fields and corresponding generation of a report definition according to an embodiment of the present invention.

FIGURE 11C is a display of a report and related structured query language (SQL) according to iterative report generation according to an embodiment of the present
15 invention.

FIGURE 11D is a display and related XML snippet illustrating the iterative addition of fields from a base view and related views according to an embodiment of the present invention.

FIGURE 11E is a display and related SQL query illustrating a report resulting from the iterative addition of fields from a base view and related views according to an embodiment of the present invention.

FIGURE 12A is a display and related XML snippet illustrating the selection
5 of fields prior to construction of a filter according to an embodiment of the present invention.

FIGURE 12B is a display illustrating the construction of a filter according to an embodiment of the present invention.

FIGURE 12C is an XML snippet illustrating the iterative generation of a query associated with the definition of a query according to an embodiment of the present
10 invention.

FIGURE 12D is a display illustrating a report and the associated SQL query based upon a filter generated according to an embodiment of the present invention.

FIGURE 13A is a display illustrating the iterative generation of an advanced filter according to an embodiment of the present invention.

FIGURE 13B is a display illustrating an exemplary filter construction window
15 according to an embodiment of the present invention.

FIGURE 13C is a display illustrating a filter construction window containing advanced logical elements according to an embodiment of the present invention.

FIGURE 13D is a display illustrating a construction window containing a Boolean AND logical grouping and associated drop area according to an embodiment of the present invention.

FIGURE 13E is display illustrating advanced filter development with
5 associated fields, operators and selection elements according to an embodiment of the present invention.

FIGURE 13F is a display illustrating a construction window containing both Boolean AND and Boolean NOT OR logical groupings and associated drop areas according to an embodiment of the present invention.

10 FIGURE 13G is a display illustrating the addition of data elements to an advanced filter generation window according to an embodiment of the present invention.

FIGURE 13H is a display illustrating a report based upon a query incorporating advanced filters according to an embodiment of the present invention.

FIGURE 13I is a tabular diagram of the XML and SQL queries generated in
15 connection with advanced filters according to an embodiment of the present invention.

FIGURE 14A is a window illustrating a method of incorporating subfilters into a report according to an embodiment of the present invention.

FIGURE 14B includes windows illustrating a means of defining subfilters according to an embodiment of the present invention.

FIGURE 14C is a window illustrating various user options for defining subfilters according to an embodiment of the present invention.

FIGURE 14D is a window illustrating definition of subfilters according to an embodiment of the present invention.

5 FIGURE 14E is a window illustrating a report incorporating subfilters according to an embodiment of the present invention.

FIGURE 14F is a tabular diagram of the XML report definition generated in connection with subfilters according to an embodiment of the present invention.

10 FIGURE 14G is a tabular diagram of SQL generated in connection with subfilters according to an embodiment of the present invention.

FIGURE 15A is a window illustrating iterative development of a drill through report according to an embodiment of the present invention.

FIGURE 15B is a window illustrating iterative development of a drill through report according to an embodiment of the present invention.

15 FIGURE 15C is a window illustrating a report incorporating drill through hot spots according to an embodiment of the present invention.

FIGURE 15D is a window illustrating a report incorporating a drill through report according to an embodiment of the present invention.

FIGURE 15E is a window illustrating field elements of a drill through report according to an embodiment of the present invention.

FIGURE 15F is a window illustrating filter elements of a drill through report according to an embodiment of the present invention.

5 FIGURE 15G is a tabular diagram of the XML generated in connection with a drill through report according to an embodiment of the present invention.

FIGURE 15H is a tabular diagram of the SQL generated in connection with a drill through report according to an embodiment of the present invention.

10 FIGURE 16A is a window illustrating iterative development of a loop back report according to an embodiment of the present invention.

FIGURE 16B is a window illustrating iterative development of a loop back report according to an embodiment of the present invention.

FIGURE 16C is a window illustrating iterative development of a loop back report according to an embodiment of the present invention.

15 FIGURE 16D is a window illustrating a report incorporating a loop back report according to an embodiment of the present invention.

FIGURE 16E is a tabular diagram of the XML generated in connection with a loop back report according to an embodiment of the present invention.

FIGURE 16F is a tabular diagram of the SQL generated in connection with a loop back report according to an embodiment of the present invention.

FIGURE 17 is a logic flow diagram illustrating a method for providing security for a report according to an embodiment of the present invention.

5

DETAILED DESCRIPTION

The present invention may be embodied in a computer database query and reporting system that groups and displays selected database data based upon base views, and the fields and relations associated with those base views. Selected database data is displayed on a display surface according to row, column, summary and group criteria chosen by a user. The display surface is typically an active window on a display device of a simple application program, but the display surface may alternately be a window of a web browser or any application program operable for displaying and manipulating data. The display surface is typically a monitor, but may alternately be a printer, flatscreen LCD display, television, and so on.

In one embodiment of the invention, the display surface includes a query construction window and a query reporting window. The query construction window includes a recursive tree structure area, a column drop area, a group drop area and a measures drop area. The recursive tree structure area is a display item used to display database views and associated fields and relations, is typically located at the left of the query construction window and is column-shaped. The group drop area is a display item used for adding fields from the recursive tree structure area to create row groupings of a report, is typically located to the right of the recursive tree structure area, and is column-shaped. The column drop area is a display item used for adding fields from the recursive tree structure area to create columns of a report, is typically located to the right of the group drop area, and is column-shaped. The measures drop area is a display item used for adding fields from the recursive tree structure area to create summary or total fields of a report, is typically located to the right of the column drop area, and is column-shaped. Alternate embodiments may use different

means of displaying the names of database fields and relations. Alternative embodiments may use more drop areas, or a single drop area. Alternate embodiments may also change the shape of the drop area display items to fit various displays; for example, the drop areas may be round, square, triangular, or a custom shape as needed, or may be located in a pull-down menu or in some other type of user interface configuration. For example, the drop areas may be located in combined windows on the display screen, or may be represented by icons or buttons rather than blank fields.

According to another embodiment of the invention, the database query and reporting system may add columns to a report. To add columns, a user selects a field from the list of fields in the recursive tree structure area and drops the field in the column drop area by initiating a drag-and-drop command, or, provided the column drop area is active, by double-clicking the desired field, or by clicking an arrow-transfer-button display item. The database query and reporting system captures this action by adding the field to the list of columns and by displaying the selected field name as a column heading name.

Similarly, in another embodiment of the invention, the database query and reporting system may add row groupings to a report. To add row groupings, a user selects a group field from the list of fields in the recursive tree structure area and drops the field in the group drop area by initiating a drag-and-drop command, or, provided the group drop area is active, by double-clicking the desired field, or by clicking an arrow-transfer-button display item. The database query and reporting system captures this action by adding the field to the list of row groupings and by displaying the selected field name as a group heading name.

According to another embodiment of the invention, the database query and reporting system may add numeric summary or aggregation measures to a report. To add measures, a user selects an aggregation or measures field from the list of fields in the recursive tree structure area and drops the field in the measures drop area by initiating a drag-
5 and-drop command, or, provided the measures drop area is active, by double-clicking the desired field, or by clicking an arrow-transfer-button display item. The database query and reporting system captures this action by adding the field to the list of measures and by displaying the selected field name as a measure heading name.

Likewise, according to yet another embodiment of the present invention, the
10 database query and reporting system may add fields from related database views to a report. Typically, the recursive tree structure area will include a list of relations, which are related tables or views. To add fields from a related database view or report, a user initiates a double-click command on a relation. The database query and reporting system will respond by replacing the previously existing list of fields and relations in the recursive tree structure
15 area with a new list based upon the selected relation and the cardinality existing between the base view and the destination view of the relation. The user may then add fields from the recursive tree structure area to the column, group and or measures drop areas, as noted above.

Once a user has selected the desired fields to be grouped and displayed in a report, in one embodiment of the invention, the user may select to view a corresponding
20 report. Typically, a user will view a corresponding report by selecting a view report button. Alternate embodiments may allow a user to review a report by changing focus on a window, or may automatically display a report after each field is added, and so on.

In the present invention, the list of fields and relations displayed in the recursive tree structure area is based upon a base view. Typically, the database query and reporting system will retrieve a list of tables and views from a database server and display them on a display surface. A user may then select one of the tables or views. Based upon the user's selection, the database query and reporting system will generate a hierarchical or entity-relationship map of all tables, views, fields and relations of the selected table or view and the tables and views related to the selected table or view.

In the present invention, access to views, fields and relations is managed by a robust security model, enabling organizations to protect sensitive data while providing broad data access. Security is enforced through the model by using access control lists and security filters. Users and groups of users, or security principals, are defined in the access control lists, and also in security filters.

Furthermore, the described features, structures, or characteristics may be combined in any suitable manner in one or more embodiments. In the following description, numerous specific details are provided, such as examples of programming, user selections, network transactions, database queries, database structures, physical structures, etc., to provide a thorough understanding of embodiments of the invention. One skilled in the relevant art will recognize, however, that the invention can be practiced without one or more of the specific details, or with other programs, user selections, database queries, database structures, physical structures, etc. In other instances, well-known structures, methods, or operations are not shown or described in detail to avoid obscuring aspects of the invention.

As used herein, a "user" refers not only to a person using the present invention, but also to a program, application, operating system, function call, or any other entity that may make use of the present invention. Thus, an operating system that manipulates or otherwise employs the present invention is classified as a user.

5 Fig. 1 and the following discussion are intended to provide a brief, general description of a suitable computing environment in which the invention may be implemented. While the invention will be described in the general context of application programs running on operating systems in a distributed computing environment where tasks are linked through a communications network, those skilled in the art will recognize that the
10 invention also may be implemented in varying types of computer environments, including desktop computers, laptops, hand-held devices, multiprocessor systems, microprocessor-based or programmable consumer electronics, minicomputers, mainframe computers, and the like. In a distributed computing environment, application programs may be located in both local and remote memory storage devices.

15 With reference to Fig. 1, according to one embodiment of the present invention, a computer system for implementing the invention includes a conventional Desktop Computer 1, an Application Server 2 and a Database Server 3. Typically, the Desktop Computer 1, the Application Server 2 and the Database Server 3 will operate in a networked environment using logical connections. Although FIG. 1 depicts a system
20 including a Desktop Computer 1, it will be appreciated by those skilled in the art that other types of computing devices such as a Laptop Computer 4, or a Personal Digital Assistant 5, may also be used.

Typically, the Desktop Computer 1 includes a Processing Unit 6, System Memory 7, and a System Bus 8 that couples the System Memory 7 to the Processing Unit 6. The System Memory 7 includes Read Only Memory (ROM) 9 and Random Access Memory (RAM) 10, and a Basic Input/Output System (BIOS) 11 that contains the basic routines that help to transfer information between elements within the Desktop Computer 1, such as during start-up, and the ROM 9. The Desktop Computer 1 further typically includes a Hard Disk Drive 12. The Hard Disk Drive 12 is connected to the System Bus 8. The Hard Disk Drive 12 and its associated computer-readable media provide nonvolatile storage for the Desktop Computer 1. Although the description of computer-readable media above refers to a hard disk, it will be appreciated by those skilled in the art that other types of storage devices and media that are readable by a computer, such as a removable magnetic disk, a CD-ROM disk, a magnetic cassette, a flash memory card, a digital video disk, Bernoulli cartridge, and the like, may also be used included in, or attached to, the Desktop Computer 1.

A number of program modules may be stored in the Hard Disk Drive 12 and the RAM 10, including an Operating System 13, one or more Application Programs 14, a Web Browser Program 15, and Program Data 16. These program modules include a Data Query And Reporting User Application (DQR Application) 100 configured for implementing an embodiment of the present invention. A user may enter commands and information into the Desktop Computer 1 through conventional input devices such as a Keyboard 17 or a pointing device such as a Mouse 18. Other input devices (not shown) may include a pen, touch-operated device, microphone, joystick, game pad, satellite dish, scanner, or the like. A Display Device 19, such as a display screen, is also connected to the System Bus 8 via an

interface. In addition to the Display Device 19, desktop computers typically include other peripheral output devices (not shown), such as speakers, scanners or printers.

Application Server 2 and a Database Server 3 may be personal computers, minicomputers or mainframe computers, or another common application platform, and may also include many or all of the elements described relative to the Desktop Computer 1.

Typically, the logical connections depicted in Fig. 1 include a Local Area Network (LAN) 22 running over an Ethernet Network Bus 23 or a Wide Area Network (WAN) 24. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet. Typically, Database Server 3 stores and manages data by means of a special set of files or folders, such as an RDBMS Data Store 21 and makes that data available to other computer programs through Application Programming Interface 27, which runs in Server Program Memory 28 of Database Server 3.

When used in a typical networking environment, the Desktop Computer 1 is connected to the LAN 22 through a Network Interface Card 25. When used in a WAN networking environment, the Desktop Computer 1 typically includes a Modem 26 or other means for establishing communications over the WAN 24, such as the Internet. The Modem 26, which may be internal or external, is connected to the System Bus 8. In a networked environment, Application Programs 20, or portions thereof, may be executed on Application Server 3 and stored in the server memory and storage devices. These application programs include a Data Query And Reporting Query Generation And Database Interface Application (Query Engine) 200 configured for implementing an embodiment of the present invention. Typically, the Query Engine 200 also includes an intermediate mapping or metadata layer

that is used when communicating with a database server. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

Fig. 2 is a block diagram illustrating the main tables, fields and the relations of a sample database, which has been derived from the Northwind database provided by Microsoft Corporation with its database server products. This modified Northwind database is used extensively in the embodiments illustrated below to show how the various embodiments of the DQR Application 100 and Query Engine 200 interact with a Data Store 21. Tables in the database are depicted in the large blocks of Fig. 2, such as a Suppliers Table 30, an Employees Table 31 and a Shippers Table 32. Fig. 2 also depicts connector lines between the tables to designate relations, such as a Relation 33 between the Employees Table 31 and the Orders Table 34. As depicted in Fig. 2, the key symbol and the infinity symbol (∞) designate the cardinality of relationships, thus the key symbol designates a “one-to” or a “to-one” relationship, and the infinity symbol designates a “many-to” or “to-many” relationship. Thus, the cardinality of the Relation 33 is expressed as one-to-many from the perspective of the Employees Table 31 in Fig. 2. As also shown in Fig. 2, the Relation 33 is linked between the EmployeeID Field 35 in the Employees Table 31 and the Employee ID Field 36 in the Orders Table 34.

Fig. 3A, Fig. 3B and Fig 3C are tables illustrating the mappings between the sample Northwind database tables, columns and relations and the views, fields and relations of an embodiment of DQR Application 100 and Query Engine 200, as used in one embodiment of the present invention. Such mappings are known by those skilled in the art as

metadata, or data describing other data. Typically, such metadata mappings are constructed by personnel familiar with a data store and the data contained therein.

In the present example metadata, a "Customer View" Table 40 depicts a mapping between the sample Northwind database described in Fig. 2 and the DQR

5 Application 100. Referring to Customer View Table 40 in Fig. 3A, a Company Name Field 41 is mapped to a CompanyName Field 42 in the Customers Table 37 of Fig. 2. Such is denoted by Balloon Number 43 in Fig. 3A. One skilled in the art will readily recognize the mappings between the metadata denoted in Fig. 3A, Fig. 3B and Fig. 3C and the tables, columns and relations of Fig. 2.

10 Figures 3D and 3E further disclose the organizational structure of the metadata. In the present invention, metadata for a database is organized in a specific manner to facilitate ad hoc data analysis thereof. In one embodiment of the present invention, metadata is organized through at least four specific software objects. Such objects have methods and properties associated with them. Table 50 of Fig. 3D describes properties
15 associated with database objects. For example, an Object Property dbUtilityTypeName D01 references a string containing the name of the object type used to access the referenced database, which could be a name readily understandable by humans or an alphanumeric reference to the database. An Object Property connectionString D02 references a string containing the location, access method and security associated with a database. One skilled in
20 the art will recognize that other property names and property types could readily be substituted for those presented in Fig. 3D. Further, one skilled in the art will also recognized

that other software conventions such as functions, structures and the like could be used instead of objects.

According to one embodiment of the present invention, instances of the objects described in Fig. 3D are implemented through use of eXtensible Markup Language 1.0 (XML). Table 60 of Fig. 3E includes an XML description of an instance of the Database object described in the Table 50 for the Northwind sample database described in Fig. 2. Referring to Fig. 3E, note that a dbUtilityTypeName Property 61 specifies that SQL Server is the access method for the Northwind database. Note also that a connectionString Property 62 indicates the Northwind database is located on the local machine and accessed through integrated security. One skilled in the art will readily recognize that different database access service providers and securities interfaces may be used.

As shown in Fig. 3E, a Table 63 includes the XML description of an instance according to the description of the Table 51 of the Customer View 40. In one embodiment of the present invention, each view described by the metadata has a corresponding XML object definition. In the Table 63, the xsi:type="view" Tag 64 specifies the object as a view object; the databaseID="1218" Tag 65 specifies a shorthand notation referencing the modified Northwind database; and the sourceTable="Customers" Tag 66 indicates that the Customer View is mapped to the Customers Table 37 in Fig. 2. The <primaryKey keyColumn="CustomerID" dataType="Text" /> Tag 67 indicates that the key field for the Customer View 40 is the CustomerID Field 38. The <defaultFields> Tag 68 enumerates the source fields displayed when the user fails to specify a field after following a relation that terminates on the Customer View 40. In the present case, the XML Tag <field

ref="northwind\Customer\Company Name" /> 69 references the Company Name Source Field 41 of Fig. 3A. The XML Tag <defaultAggregateFields> 70 enumerates the source fields containing numeric values associated with the Customer View 40, which are available for providing numeric summaries of data contained in a report. In the present embodiment, 5 the XML Tag <field ref="1228" type="aggregate" /> 71 references the Customers Aggregation Field 44 of Fig. 3A.

Table 72 of Fig. 3E provides an XML description of the Address Field 45 of the Customer View 40 of Fig. 3A and the Customer View XML Object 63 in Fig. 3E. In one embodiment of the invention, each source field to be exposed for a view is similarly defined. 10 The xsi:type="savedSourceField" XML Tag 73 identifies an Address Object 72 as a data or source field. The sourceColumn="Address" XML Tag 74 identifies the Address Field 39 as the data source for the Address Object 72.

As shown in Fig. 3E, an Orders relation Object Table 76 is an instance of a Relation object conforming to the Table 53, which provides an XML description of the 15 Orders Relation 46 of Fig. 3A. According to one embodiment of the invention, each relation is similarly defined. Referring to the Orders Relation Object Table 76 of Fig. 3E, an xsi:type="relation" Tag 77 defines the object as a relation object. The relation definition also includes a ViewID="northwind\Order" Property 78, which in the present embodiment indicates that following a relation from the Customer View 40 to the Order View 47 will 20 expose the fields and relations associated with the Order View 47. A reverseID="northwind\Order\Customer" Property 79 indicates, should the Orders Relation 46 be followed, that the path back to the Customer View 40 will occur through use of the

Customer Relation 48. A relationType="OneToMany" Property 81 indicates that the relation from the Customer View 40 to the Order View 47 is one-to-many. The join type and the join keys for the Orders relation Object Table 76 are specified by a joinType="LeftOuterJoin" property 80 and the <joinKey sourceColumn="CustomerID" destColumn="CustomerID" dataType="Text" /> XML Tag 82, respectively. In this case, because the relationship is identified as a one-to-many relation, the join is specified as a left outer join. A left outer join of the Customer View 40 and the Order View 47 will include all records from the Customers Table 37 and the corresponding records in the Orders Table 34 where the CustomerID 38 and the CustomerID 38A are equal.

A Table 83 of Fig. 3E includes an XML description of the Customer Relation 48 of the Order View 47 of Fig. 3B. The Table 83 represents the reverse path associated with the Orders Relation 46. In this case, a toViewID="northwind\Customer" Property 84 points to the Customer View 63, a reverseID="northwind\Customer\Orders" Property 85 points to the Orders Relation Object Table 76, a relationType="ManyToOne" Property 86 indicates that the relation is many-to-one, and a joinType="InnerJoin" Property 77 indicates that the join is an inner join. An inner join will include records from both the Orders Table 34 and the Customers Table 37 where the values of the join keys specified by the <joinKey sourceColumn="CustomerID" destColumn="CustomerID" dataType="Text" /> XML Tag 88 are equal.

One embodiment of the present invention provides a user the means to iteratively generate queries. Fig. 4 depicts a main or initial display Window 92 of an

embodiment of the DQR Application 100. From this Window 92, a user of the DQR Application 100 may select a New Button 90 to create a new report.

Fig. 5 depicts a display Window 94 according to one embodiment of the DQR Application 100 that is useful for guiding a user through the process of creating a report. One skilled in the art will appreciate that various other interfaces may be used to facilitate creation of a report, including a menu-drive interface, a programmatic interface, a verbal interface, etc. In the embodiment shown, a user may select a Detail Report Radio Button 110 to create a new detail report. A detail report in the depicted embodiment is a list based upon one or more source fields of a database view. In this embodiment, a user may also select a Crosstab Report Radio Button 111 to create a tabulated report based upon the intersection of two source fields that bear a many-to-many relationship to each other. A user may also select a Based Upon Existing Template Radio Button 112 to create a report based upon a previously saved report or template. Once a user has selected a report type, a user may select a Next Button 113 to proceed to the next step in creating a report. A user may also select a Cancel Button 114 to stop building a report, or may select a Back Button 115 to return to the Window 92.

Figures 6A and 6B depict display Windows 96A and 96B, respectively, of an embodiment of the DQR Application 100 that allows a user to select a database view as the base view for use in a list type report. The user may select a database known to the DQR Application 100 and the Query Engine 200, as described by metadata associated with the database (an example of which is provided in Figures 3A-3E) by selecting one of the databases included in a Look In Drop Down Box 120. In the depicted embodiment, once a

user has selected a database, the views associated with the database as described by the metadata are displayed in a View List 121 of Fig. 6B. In the present example, the views described in Figures 3A – 3C appear in the View List 121.

In the depicted embodiment, views may be organized into subfolders, such as a Lookup Folder 122. This facility is provided for databases having a large number of defined views. In the depicted embodiment of the present invention, the base view is set by selecting a view from a List 123 and either selecting a Finish Button 124 or double-clicking on the selected view. This base view, in conjunction with the associated metadata described in Figures 3A-3C, as exposed by the Query Engine 200 and the DQR Application 100, is used in the invention to provide an entity-relationship map of the database relative to the selected base view.

The entity-relationship map of the present invention may be exposed through a database entity-relationship grouping and display system and according to rules of the invention enforcing proper display and element selection and iterative query generation. Fig. 7 depicts a display Window 130 of the grouping and display system according to one embodiment of the DQR Application 100. Those skilled in the art will recognize that the Window 130 is a conventional window of a modern desktop application. However, those skilled in the art will also recognize that other conventional and non-conventional display means, screens and windows could be used. In the depicted embodiment, a Menu Bar 131 contains several menu items, including a View Menu Item 132, which in one embodiment of the invention shows the iterative nature of queries generated. A Toolbar 133 contains a Filters Button 135, a Sort Button 136, an Options Button 137 and a View Report Button 138.

The Toolbar 133 and related buttons are used in this embodiment of the invention to display certain windows and build iterative queries.

A Group By Box 139, a Details Box 140 and a Measures Box 141 are standard label boxes. The boxes below the Group By, Details and Measures areas of the display window, numbered respectively 130, 131 and 132, are drop box areas where a user may drag or locate fields when building queries. Note that in the depicted embodiment, a Details Drop Box Area 143 has been selected by default, as denoted by the darker gray colored background surrounding the Details Label 140. With such selection, a user may select from a Selection Area 146 one or more fields from a Field Group 147 that pertains to the Employee View 49, which is the view based upon the Employees Table 31 of the Fig. 2, as shown in a Look In Drop Down Box 145 of Fig. 7. In the depicted embodiment and present example, one of the fields directly associated with the Employee View 49 is a Full Name Source Field 148, which has been selected according to the present invention as shown in the Details Drop Area 143. In the present example, a user may also follow relations between the Employee View 49 and the other views described in Figures 3A-3C, including an Employee Territories Relation 150, an Orders Relation 151, a Reports To Relation 152 and a Subordinates Relation 153. In the present example, the Reports To Relation 152 is visually designated as a to-one relation, as denoted by a superscripted "1" 154.

Fig. 8 depicts a display Window 158 of an embodiment of the DQR Application 100 that illustrates selection of a Group By Drop Area 160. When the Group By Drop Area 160 is selected, the DQR Application 100 displays only those fields of the Employee View 49 as are designated in a Look In Drop Down Box 161 that can be used for

grouping. Each such field will have been designated previously as a field available for grouping by setting an allowAsGroupField Property F05 (see Fig. 3D) to true. In the present example, the fields highlighted by a Balloon 162 may be selected and the relations highlighted by a Balloon 163 may be followed.

5 Fig. 9 depicts a display Window 186 of an embodiment of the DQR Application 100 illustrating selection of a Measures Drop Area 170. In the depicted embodiment, the Measures Drop Area 170 provides a means to incorporate numeric summaries or totals into a report. When the Measures Drop Box 170 is selected, only those fields pertaining to the selected view, in the present example the Employee View 49, as
10 designated in a Look in Drop Box 171, are available for selection. In the present example, the total number of employees in the Northwind database of Fig. 2 can be added to the Measures Drop Box 170. A user may also choose to follow one of the relations identified by a Balloon 172. In the present embodiment, a user chooses to follow a relation by double-clicking on the relation name. Should a user drag a relation name to the Measures Drop Area 170 or left-
15 click on a relation name followed by clicking an Add Button 173 while the Measures Drop Area 170 is the default drop area, the DQR Application 100 will add the fields identified by a defaultAggregateFields Property V06 (See Fig. 3D) for the view to the Measures Drop Area 170.

 Fig. 10A displays a flowchart detailing the steps of operation of the method of
20 selecting a base view, using that base view to constrain data selection, and creating reports. In step 180, a relational abstraction of a data store is created. Steps 181, 182 and 183 describe the steps of creating the relational abstraction. In Step 181, views of the data to be available

for access from the data store are created. In Step 182, fields to be associated with such views are defined. In Step 183, relations between views are defined. In one embodiment of the present invention, a metadata layer that references the Northwind sample database is created, as depicted in Figures 3A through 3E and as described above. One skilled in the art will

5 readily recognize that a similar abstraction may be created for data stores of all types, including data stores that are not based upon relational database methodologies.

Once a relational abstraction of a data store has been created, reports may be generated. In Step 184, a view is selected. This is the base view of a report. Based upon that selected base view, fields may be selected in Step 185. Such fields may be either directly
10 associated with the base view, or they may be associated with other views defined in Step 182. In Step 186, a relation path from the base view selected in Step 184 to a destination view is determined by following a series of relations defined in Step 183. If no relations are followed, the relation path is empty. A relation path indicates whether the cardinality of a related view relative to the base view is to-one or to-many.

15 In Step 187, a decision is made based upon the relation path determined in Step 186. If the cardinality of the relation path is to-one, a scalar report field may be created by selecting a field in the destination view, or creating an expression that references one or more fields, as shown in Step 191. If the cardinality of the relation path is to-many, an aggregate report field may be created by selecting or creating one or more fields whose base
20 view is the last view of the relation path, as shown in Step 188, and specifying a function or expression for aggregating the values of the nested fields, as shown in Step 189. Such functions or expressions may be of various types. For example, a function may simply sum

the values of a nested field. In the Northwind sample database, such a function might sum the shipping costs of all orders for a particular customer. According to one embodiment of the invention, in Step 192, a scalar field or expression from step 191, or a function or expression from Step 189 is added to a report definition. The field selection process as outlined in Steps 5 185 through 192 is repeated recursively until the report definition meets the requirements of a user, whether that user is a person or a process.

Fig. 10B displays a flowchart detailing the steps of operation of a database entity-relationship grouping and display system according to one embodiment of the invention. In Step 201 an embodiment of the DQR Application 100 running on the Desktop 10 Computer 1 requests a list of available databases and the list of metadata views identified in Figures 3A, 3B and 3C, from the Query Engine 200 running on the Application Server 2. The Query Engine 200 responds with the names of available databases and views, including the Northwind example database described in Fig. 2, and displays them through the Window 96A shown in Fig. 6.

15 In step 202, the DQR Application 100 displays the name of the Northwind database described in Fig. 2 and a list of the other available databases in the Look In Drop Down Box 120 on the Display Device 19. Step 202 also displays the metadata views described in Figures 3A, 3B and 3C in the View List Box 121 through the Window 96B shown in Fig. 6A.

20 In Step 203, a user sequentially selects the Employee View 49 from the View Group List 123 and the Finish Button 124, which causes the YES branch of Step 203 to be

followed. If a user does not select a view, the NO branch of Step 203 is followed and the DQR Application 100 continues to display the Window 96B of Fig. 6A.

In Step 204, the DQR Application 100 requests the fields and relations listed in the Employee View 49 from the Query Engine 200. In Step 205, the DQR Application 100
5 then displays the list of fields and relations of the Employee View 49 on the Display Device 19, displaying the Window 130 described in Fig. 7.

In Step 206, the DQR Application 100 awaits user input in the form of selecting fields, such as those highlighted by the Balloon 147, or relations, such as those highlighted by Balloon 149, in Fig. 7. If the View Report Button 138 is selected, the YES
10 branch of step 206 is followed to Step 218. If no fields have been added, the NO branch of Step 218 is followed to Step 219, an error is displayed directing the user to select at least one field, and the DQR Application 100 continues to display the list of fields and relations of the selected view, such as the Employee View 49. If the user selects a field or relation, the DQR Application 100 proceeds to Step 208.

15 In Step 208, the DQR Application 100 monitors detail field selections. If a user does not select a detail field, the DQR Application 100 continues through the NO branch to Step 209. If a user selects a detail field, the DQR Application 100 proceeds through the YES branch to Step 212. In Step 212, the DQR Application 100 adds the name of the selected field to the Details Drop Box Area 143, and continues to Step 205 to display the
20 fields and relations associated with the selected view.

In Step 209, the DQR Application 100 monitors group field selections. If a user selects a group field, the DQR Application 100 proceeds through the YES branch to

Step 213. If a user does not select a group field, the DQR Application 100 continues through the NO branch to step 210.

In Step 210, the DQR Application 100 monitors measure field selections. If a user selects a measure field, the DQR Application 100 proceeds through the YES branch to

5 Step 214. If a user does not select a measure field, the DQR Application 100 continues through the NO branch to Step 211.

In Step 211, the DQR Application 100 monitors the selection of relations. If a user selects a relation, the DQR Application 100 proceeds through the YES branch to Step

215. If a user does not select a relation, the DQR Application 100 continues through the NO
10 branch to Step 205.

In Step 215, if the cardinality of the relation path ending with the selected relation is to-one, the DQR Application 100 follows the NO branch of Step 215 to Step 216.

In Step 216, the DQR Application 100 retrieves the fields and relations associated with the followed relation and processing passes to Step 205. If the cardinality is to-many, the YES
15 branch of Step 215 is followed to Step 217 where the DQR Application 100 limits retrieval of the fields associated with destination view to those fields that have a Field Type Property F07 (see Fig. 3D) set to "aggregate" and then processing is passed to Step 205. In this manner, the cardinality of the destination view relative to the base view constrains field
20 selection. If the cardinality is to-many, only aggregated values associated with the destination view may be returned, thereby ensuring that each row returned by the DQR Application 100 represents exactly one row in the base view selected for the report.

In Step 218, if at least one field has been added to the report, the YES branch is followed to Step 220, where the DQR Application 100 verifies and generates a suitable database query and displays the report on the Display Device 19. In Step 221, if the Fields Button 134 is selected, the YES branch is followed to step 205. Otherwise the process
5 terminates.

Figures 11A through 11J depict the iterative manner of query generation according to one embodiment of the present invention. In each case, a change of selected fields iteratively and automatically changes an XML file, which file contains one embodiment of a self-contained report definition of the present invention. Fig. 11A depicts a
10 Window 250 and an XML Report Definition Table 251 according to one embodiment of the invention. The Window 250 illustrates the selection of the Employee View 49 in a Look In Box 252. The absence of any fields in a Group By Drop Area 253, a Details Drop Area 254 and a Measures Drop Area 255 of the Window 250 indicates that in the present example no fields have been selected. The XML Definition Table 251 was automatically generated by the
15 DQR Application 100 upon selection of the Employee View 49. A. baseViewID="Employee" XML Tag 252 specifies that the starting view of the report is the Employee View 49.

Fig. 11B depicts a Window 260 and an XML Report Definition Table 261. The Window 260 is an iteration of the Window 250. The XML Report Definition Table 261
20 is an iteration of the XML Report Definition Table 251. In the Window 260, the Full Name, Title and Hire Date fields of the Employee View 49 have been added to the Details Drop

Area 262. A <detailFields> XML Tag 263 indicates the addition of the Full Name, Title and Hire Date fields of Employee View 49.

Fig. 11C includes a Window 270 and an SQL Query Table 271. The Window 270 results from user selection of a View Report Button 272 and the selections described

5 with reference to Fig. 11B. The Window 270 includes a Report Table 273. The Report Table 273 includes a Header Row 274 and Several Data Rows 275. The table 273 is comprised of three columns, Columns 276, 277 and 278. Each of these Columns corresponds to the detail fields inserted in the <detailFields> XML Tag 263 as recorded in the XML Report Definition Table 261. The SQL Query Table 271 was derived from the XML Report Definition Table
10 261. One skilled in the art will readily recognize the syntax and structure of the SQL statements included in the SQL Query Table 271.

Fig. 11D includes a Window 280 and an XML Report Definition Table 281.

The Window 280 is an iteration of the Window 260. The XML Report Definition Table 281 is the corresponding iteration of the XML Report Definition Table 261. Several changes have

15 been made between the Window 260 and the Window 280, including that the Hire Date field of the Employee View 49 has been removed from a Details Drop Area 284, the default field for the Country relation of the Employee View 49 has been added to a Group By Drop Area 283, and the Orders relation of the Employee View 49 has been followed and the Orders Aggregation Field 286 of the Order View 47 has been added to a Measures Drop Area 285.

20 These changes have caused several iterative changes to the XML Report Definition Table 281, including removal of the Hire Date field from an XML Detail Fields 287, addition of a new XML section, noted by a <groupFields> XML Tag 288 and addition of the default field

of the Country relation of the Employee View 49 as a group, addition of a new XML section, noted by a <measureFields> XML Tag 289, and addition of the Orders Aggregation Field 286 of the Order View 47.

Fig. 11E includes a Window 290 and an SQL Query Table 291. The Window 290 results from conversion of the XML Report Definition Table 281 into SQL and execution of the query contained in the SQL Query Table 291. The resulting report depicted in the Window 290 includes a Full Name Column 293 and a Title Column 294. The Window 290 also illustrates a Country Group 292 and a Summary 295, which result from the Orders aggregation field of the Order View 47. One skilled in the art will recognize the SQL query included in the SQL Query Table 291.

Figures 12A through 12D show windows and reports incorporating the use of filters for the present invention. Fig. 12A illustrates a Window 300 and an XML Report Definition Table 301. The Window 300 is an embodiment of the present invention based upon selecting the Order View 47. The Order Date field of the Order View 47 and the Country Name field of the Ship Country relation have been included in a Details Drop Area 303. The Orders aggregation field of the Order View 47 has been included in a Measures Drop Area 306. Based upon following the Customer Relation 48 of the Order View 47, the default field Company Name of the Customer View 40 has been included in a Group By Drop Area 305. The XML Report Definition Table 301 includes the iteratively created XML definition corresponding to the foregoing selections.

Fig. 12B depicts a Window 310, which discloses how the entity-relationship or hierarchical map may be used in the present invention to construct a report filter. The

Window 310 results from user selection of a Filters Button 311. Fields from a List 312 may be selected by dragging and dropping field names to a Panel 313. Likewise, fields from the List 312 may be selected by double-clicking or by selecting a Button 314. Relations from a List 315 may also be followed, leading to other fields that may be selected similarly. As presented, the Panel 313 illustrates the selection of an Order Date Field 316 and a Ship Country Field 320. The Order Date Field 316 is connected to a Drop Down Box 318 by a Between Operator 317. Several other operators are available for filters, including "Equals," "Greater than," "Less than," "Greater than or equal to," "Less than or equal to," "In" and "Not." Coupled with the Between Operator 317, Drop Down Boxes 318 and 319 indicate that the current query will return only records with the Order Date 316 between January 1, 1996 and December 31, 1996. The Ship Country Field 320 is connected to a Country List 322 by an In Operator 321, disclosing that the current query will return only records having a value of Ship Country equal to the list of countries in the Country List 322. An XML Report Definition Table 324 depicted in Fig. 12C results iteratively from the selections described with reference to the windows shown in Figures 12A and 12B. Fig. 12D contains a Report Window 326 and an SQL Query Table 328, both of which result from execution of the query disclosed in Fig. 12C.

Figures 13A through 13I show windows and queries for advanced filter options according to the present invention. Heretofore, users have been subjected to extremely complex scripting languages or notational Boolean logic elements to build advanced queries. The present invention provides a simplified filter display and creation system.

Fig. 13A illustrates a Window 330 according to one embodiment of the present invention. According to a Look In Drop Down box 331 of window 330, the Order View 47 has been selected as the base view. The Ship Country field of the Order View 47 has been selected as a group field and is displayed in a Group By Drop Area 332. The Order Date and Order Total fields of the Order View 47 have been selected as columns and are displayed in a Details Drop Area 333.

Fig. 13B illustrates a Window 340, an iteration of the Window 300 caused by selection of a Filters Button 31. The Window 340 represents an embodiment of the main filter display and generation window of the present invention. The Window 340 includes a Display Box 342, which contains a listing of fields and relations according to the Order View 47, which is the base view. The Window 340 also includes Blank Panel 343, which is a panel used for displaying filters upon creation and manipulation. The Window 340 also includes an Expand Button 344, which is used to expand filter options based upon logical groups, as denoted by a Display Box 345.

Fig. 13C illustrates a Window 350, which results from selection of the Expand Button 344 of Window 340. The Window 350 includes a Logical Groups Box 351, which contains four logical elements, numbered 352 through 355. A Logical Element (Any of (OR)) 352 is useful for building filters according to a Boolean OR. A Logical Element (All of (AND)) 353 is useful for building filters according to a Boolean AND. A Logical Element (None of (NOT OR)) 354 is useful for building filters according to a Boolean NOR. A Logical Element (Not all of (NOT AND)) 355 is useful for building filters according to a Boolean NAND. One skilled in the art will recognize that logical elements could be incorporated into

Window 350 to cover the other cases of Boolean logic. Such logical elements can be selected by double-clicking or dragging and dropping the logical element on a Blank Panel 357.

Alternatively, a logical element may be selected by left-clicking on the logical element and a Selection Button 356.

5 As depicted in Fig. 13D. A Window 360 illustrates the addition of the All of (AND) Logical Element 353 from Window 350. Upon selection, a Text Box 361 is displayed on a Panel 362. In addition, an Indented Drop Area 363 appears, indicating that fields may be dropped there to build a filter group. A Drop Area 364 also appears, indicating that additional filter groups or items may be added at the top level.

10 A Window 370, depicted in Fig. 13E, illustrates the addition of the Order Date and Shipper fields of the Order View 47 to the Window 360. As fields are added to the Indented Drop Area 363, fields, operators and applicable window elements displace the Drop Box, and the DQR Application 100 displays another Drop Box 371, indicating that additional fields may be added. In the Window 370 the Order Date field 372 is followed by a Greater
15 Than Operator 373, which is aligned with a Calendar Drop Down Box 374. Selecting the Calendar Drop Down Box 374 causes the DQR Application 100 to display a standard monthly calendar for date selection. Likewise, a Shipper Field 375 is followed by an Equals Operator 376 that is aligned with a Drop Down Box 377. Selecting the Drop Down Box 377 causes the DQR Application 100 to retrieve and display the list of Shippers from the
20 Shippers Database Table 32.

 As depicted in Fig. 13F, a Window 380 shows the result of the addition of an additional logical element to the Window 370. In this case, a None of (NOT OR) Logical

Element 381 has been selected. Upon selection, text 383 is added to a Panel 382 and a Drop Area 384 is displayed.

Upon the addition of a field to the None of (NOT OR) Logical Element 381 of the Window 380, a Window 390 is generated as depicted in Fig. 13G. According to the present embodiment of the invention, a Ship Country Field 391 appears below the None of (NOT OR) Logical Element 383. The Ship Country Field 391 is also followed by an Equals Operator 392, which in turn is followed by a Drop Down Box 393. The Drop down box is linked to the Country table of the Northwind database shown in Fig. 2. In the present case, USA appears selected in the Drop Down Box 393. As depicted in the Window 390, the constructed query would return all records in the Orders Table 34 of Fig. 2 that were shipped outside of the United States using Federal Shipping after 1996.

A Window 395 depicted in Fig. 13H illustrates the results of running the query created in Figures 13A through 13G. Note that the filter constructed in Figures 13A through 13G is displayed in the Window area 396.

Fig. 13I illustrates the XML report definitions and SQL query generated by the DQR Application 100 in response to the steps depicted and explained with reference to Figures 13A through 13G. Section 398 has been added to the XML report definition to define the report filter.

Figures 14A through 14F show windows and queries for building subfilters according to one embodiment of the present invention. According to the present invention, subfilters provide a means of filtering the rows that contribute to an aggregated value.

Fig. 14A illustrates a Window 400 of an embodiment of the DQR Application 100. In the Window 400, a report has been set up that is based on the Customer View 40, as noted in a Window Title Block 401, which will be grouped on a Country Field 402 and will display columns for a Company Name 403 and Orders 403. Note that Orders 403 is based upon following the Orders Relation 46 of the Customer View 40, which includes by default the Orders aggregation field of the Order View 47. If a report were run based upon the situation shown in the Window 400, all customers contained in the Customers Table 37 would be returned, and each customer would occupy a single row in the report. The resulting report would include 3 columns, one for the group by value of Country, and two detail columns, one containing the results of retrieving the CompanyName field of the Customers Table 37 and one containing the count of all orders contained in the Orders Table 34 associated with each customer. Note that an Icon 404 indicates that properties may be set for the associated field, in this case the Orders aggregation field.

Fig. 14B depicts a Window 405 that results from selecting the Icon 404 of the Window 400. The Window 405 includes several elements, including a General Tab 406 and a Filters Tab 407. The General Tab 406 displays a Name Label 408 that is followed by a Text Box 409. The text Box 409 contains the name of the column as it will appear in a report and can be modified by the user. In the present case, the Text Box 409 contains "Orders." The General Tab 406 also contains a Description Label 410 and a Text Box 411. The Text Box 411 contains text that typically gives a description of what the associated field contains. In the embodiment depicted in Fig. 14B, the Text Box 411 describes that the Orders aggregation field of the Order View 47 contains the number or count of orders associated with a customer. The Window 405 also includes a Save As Button 412 and a Close Button 413.

Selection of the Save As Button 412 causes the DQR Application 100 to save the current field as a persistent object that can be reused in this or other reports. The Close Button 413 causes the Window 405 to close and return to the Window 400. When the Close Button 413 is selected any changes made to field properties are maintained for the current report only.

5 A Window 415 shown in Fig. 14B displays the window elements associated with selecting the Filters Tab 407 of the Window 405. The Window 415 displays a Text Label 416, which illustrates the currently selected Orders field, a Filter Operator 417, which illustrates an In Hotlink 417 as the filter operator, and a Text Label 418, which indicates that it can be selected as a means of building a subfilter. A subfilter can be constructed according to the present window because the Orders field is an aggregation field of the Orders Relation 10 46, which relation defines a one-to-many relationship from the Customer View 40 to the Order View 47. Subfilters are meaningful only where a relationship is to-many and an aggregate field is thus required.

 A Window 420 shown in Fig. 14C depicts a subfilter wizard useful for building subfilters. This Window 420 results from selecting the Text Label 418 shown on 15 the Window 415. The Window 420 includes several window elements. Window Element 421 is a radio box that permits the user to create a new subfilter. Window Element 422 is a list box that contains any subfilters previously defined during the report development session. Window Element 423 is a radio button that permits a user to select a previously defined and 20 saved Order subfilter. Window Elements 424 and 425 permit a user to use a previously saved Order list or report as a subfilter. Window Element 426 is a button that permits the user to proceed to the next step in creating a subfilter.

A Window 430 shown in Fig. 14D illustrates the results of first selecting the Window Element 421 and then the Window Element 426 of the Window 420 (see Fig. 14C) and then creating a subfilter. The Window 430 includes several familiar window elements. The Window Element 430 includes a List 431 that includes the fields and relations associated with the Order View 47. A Window Element 432 includes the logical elements discussed in connection with Fig. 13C. A Window Panel 433 includes a Label 434 that results from selecting a Logical Group 435. An Order Date Label 436 representing the Order Date field associated with the Order View 47 has been added to the Logical Filter 434. The Order Date Label 436 is connected to Calendar Drop Down Boxes 437 and 438 by a Between Operator 439. A Shipper Label 440 has been added to the Logical Group 434 and is connected to a Drop Down Box 441 by an Equals Operator 442. As demonstrated by comparison of Figures 14A and 14D, when a subfilter has been created and incorporated into a report, an embodiment of the present invention will return a three column report of all customers, including a column showing the number of orders each customer placed in 1997 that were shipped using United Package.

Window 450 shown in Fig. 14E illustrates the results of running the query generated by the selections made in the process described with reference to in Figures 14A through 14D. Figures 14F and 14G disclose the resulting XML report definition and SQL query generated by the DQR Application 100 in creating the report displayed on the Window 450.

The process of an automatic drill through according to one embodiment of the present invention is depicted in Figures 15A through 15H. Drill through is displaying a report

containing the row or rows represented by a field value in the results of another report. The base view concept of the present invention allows drill through reports to be created automatically, using the relation path of the field and any applicable filters, including the report filter, group value filters, and aggregate subfilters as necessary.

5 Fig. 15A illustrates a Window 460 of DQR Application 100. The Window 460 is based on the Employee View 49 as the initial base view, as shown in a Look In Box 461. A Window 470 shown in Fig. 15B discloses an iterative change in the Window 460, wherein the Country relation of the Employee View 49 has been added to Group By A Drop Area 471, the Full Name field of the Employee View 49 has been added to a Details Drop Box 472
10 and the Orders aggregation field of the Order View 47 has been added to a Measures Drop Area 473.

 Executing the query resulting from selecting a View Report Button 474 shown on the Window 470 results in a Window 480, depicted in Fig. 15C. The report displayed in the Window 480 has three columns, a Country Column 481, a Full Name Column 482 and an
15 Orders Column 483. The Country Column 481 was generated because the Country relation of the Employee View 49 was added to the Group By Drop Area 471. The Full Name Column 482 resulted from adding the Full Name field of the Employee View 49 to the Details Drop Box 472. The Orders Column 483 resulted from placing the Orders aggregation field of the Order View 47 in the Measures Drop Area 473. Because the Orders Column 483 is an
20 aggregation field, each of the numbers represented in the Orders Column 483 represents the count of orders from the Orders Table 34 for each employee in the Employees Table 31. In the present example, each of the numbers in the Orders Column 484 can be left-clicked with

a mouse to access a hot spot or link associated with the number. This hot spot capability is traditional in many documents and is available in most web browsers. In the present case, the DQR Application 100 provides such a hot spot or link capability based upon the allowDrill Field Object Property F06 according to the Table 52 of Fig. 3D. In the present case, the
5 allowDrill property of the Orders aggregation field of the Order View 47 is true and accordingly any of the numbers in the Orders Column 483 can be selected or “drilled on” as it is typically described.

Selecting the Hot Spot 484 causes the DQR Application 100 to generate a Window 490, as depicted in Fig. 15D. The Window 490 is a report, as denoted by the
10 highlighting of a View Report Button 491. However, the Window 490 is not a report based upon the Employee View 49. Instead, it is a report based upon the Order View 47. A Tab 492, which includes the text “Order” reveals this fact, as does a Label 499. Columns 493 through 498 also evidence this fact. The fields associated with Columns 494, 495, 496 and 498 are all from the Order View 47. However, a label 500 reveals that the query that
15 generated the Window 490 includes a filter.

A Window 510 shown in Fig. 15E results from selecting a Fields Button 501 of the Window 490. The Window 510 is a window associated with the report shown in the Window 490. This association is highlighted by a Tab 511, which continues to display “Orders.” The Window 510 illustrates the group fields, detail fields and measures fields in
20 generating the Window 490. Several fields have been added to a Details Drop Area 512. Likewise, several fields have been added to a Measures Drop Area 513.

The filters associated with the Windows 490 and 510 can be seen by selecting a Filters button 514. Selection of the Filters Button 514 generates a Window 520 as depicted in Fig. 15F. The Window 520 illustrates one filter, a Filter 521. In this embodiment of the invention, the Filter 521 comprises three elements, a Label 522 that includes the word
5 "Employee," a Label 523 that includes the "In" operator, and a Label 524 that includes the name "Buchanan, Steven." The Label 524 refers to the name associated with the Hot Spot or Link 484 of Fig. 15C. Thus, the drill through operation produced a report based upon the Order View 47, but filtered by the parameters of Employee View 49 that were selected prior to selecting the Hot Spot 484. Figures 15G and 15H disclose the XML report definition and
10 SQL query that produced the Windows 490, 510 and 520.

Figures 16A through 16F depict the windows and processes of a loop back operation according to one embodiment of the present invention. Loop back is the concept of following relations of a database from one view through one or more other views back to the original view. According to the present invention, loop back allows non-technical users to
15 build complex queries by merely following the relations exposed by views defined in the meta data.

Fig. 16A depicts a Window 550 of the DQR Application 100. The Window 550 illustrates the start of a loop back query. The query is based upon the Order View 47, as noted by a Title Bar 551 and a Drop Down Box 552. Two fields from the Order View 47, the
20 Order Date field and the Order Total field, have been dropped on a Details Drop Area 553. One field from the Customer View 40, the Company Name field, has been dropped on the Details Drop Area 553 as well. A relation 554 depicts the Employee Relation 52 of the

Order View 47 that will be followed to build a loop back query. With respect to the Employee Relation 52, the Window 550 provides visual feedback through a superscripted "1" noted by a Balloon 555, that the relation is to-one.

Fig. 16B depicts a Window 560 that illustrates the fields and relations associated with following the Relation 554 to the Employee View 49. The Window 560 also illustrates that the Full Name Field 53 has been dropped on a Details Drop Area 561. The Window 560 also displays the Orders Relation 54 of the Employee View 49, which will be followed back to the Order View 47 to create a loop back.

Fig. 16C depicts a Window 570 that illustrates the fields and relations associated with following the Relation 562 to the Order View 47. However, the Window 570 illustrates fields of the Aggregate Field Type F07 (see Fig. 3D) only. Thus, the Window 570 illustrates the constraining nature of relationships in the present invention. The loop back query was started with the Order View 47, meaning each row returned represents a single row in the Orders Table 34. The Employee Relation 52 is a many-to-one relationship, meaning that a query incorporating related rows from the Employee View 49 will only return one Employee row for each corresponding row of the Order View 47. Accordingly, all of the source fields of the Employee View 49 could be selected without aggregation. However, because the relationship through the Orders Relation 54 back to the Order View 47 is one-to-many, only aggregate fields are available. The Window 570 also illustrates that the Orders and Avg Order Total aggregate fields were added to a Details Drop Area 571.

Fig. 16D shows a Window that results from selecting a View Report Button 581 and illustrates a report based upon the loop back query constructed using the DQR

Application 100, as discussed with reference to Figures 16A, 16B and 16C. A Balloon 582 illustrates that the number of records returned by the loop back query is 830. This is the exact number of rows contained in the Orders Table 34. The Window 580 also illustrates a multicolumn table containing the data resulting from the loop back query. Columns 583 and 585 show data from the Order Date and Order Total fields selected according to Fig. 16A from the initial or base view, the Order View 47. Column 584 includes data from the Company Name field of the Customer View 40. Column 586 includes data from Full Name field of the Employee View 49. Columns 587 and 588 include data from Orders and Avg Order Total aggregate fields, respectively, of the Order View 47 according to selecting the Orders Relation 54 of the Employee View 49, as depicted in Fig. 16C. Figures 16E and 16F show the XML report definition and SQL query generated by the DQR Application 100 in building the loop back query. The loop back query presented herein illustrates how a user may construct a complex query in a simple manner.

Fig. 17 displays a flowchart depicting the steps of the method of providing security for a report. In Step 600, security principals are defined. Security principals include users or groups of users. In Step 601, access control entries are made to a file or other access control list maintenance system or facility. The access control entries define elements of data in a data store to be protected and whether security principals are to be given or denied access to the data elements. Data elements may be course grained, meaning that access control can be implemented on data elements that contain or include other data elements, such as databases and views.

In Step 602, security filters are specified for views that require protection according to organizational needs. In Step 604, if a security filter is defined on a view, at least one row pertaining to the view is included in the security filter and is enabled for access. In Step 605, each data element to be protected in a view is included in an access control entry
5 and security principals to be denied or granted access to the data element are included in the access control entry. In Step 603, security is enforced on a report.

In Step 606, a security principal desiring to create or execute a report is authenticated. This authentication will typically be accomplished through user authentication means characteristic of computer systems. In Step 607, assuming that a security principal has
10 been authenticated and attempts to create or run a report, the security filter, as described in the Step 602, will be applied to the report by comparing the authenticated security principal with the access control entry contained in the security filter for each data element. If the authenticated security principal does not match an access control entry contained in the security filter, the authenticated security principal will be given access the corresponding
15 data.

While specific embodiments and applications of the present invention have been illustrated and described, it is to be understood that the invention is not limited to the precise configuration and components disclosed herein. Various modifications, changes, and variations which will be apparent to those skilled in the art may be made in the arrangement,
20 operation, and details of the methods and systems of the present invention disclosed herein without departing from the spirit and scope of the invention as defined by the following claims.

It will be obvious to those of skill in the art that the invention described in this specification and depicted in the FIGURES may be modified to produce different embodiments of the present invention. Thus, the present invention has several advantages over the prior art without sacrificing any of the advantages of the prior art. Although two
5 embodiments of the invention have been illustrated and described, various modifications and changes may be made by those skilled in the art without departing from the spirit and scope of the invention.